Deep Learning Important Questions PDF

The most important Deep Learning questions are the perceptron and multilayer perceptron, activation functions, backpropagation, the vanishing gradient problem, gradient descent variants and optimisers such as Adam, regularisation with dropout and batch normalisation, CNN layers with output size calculations, popular CNN architectures, RNNs, LSTM and GRU, autoencoders, GANs, transfer learning and an introduction to transformers and attention.

Key takeaways

- Backpropagation, CNN output size calculations and LSTM gates are the most common long questions.
- Learn the formulas for common activation functions and their derivatives.
- Draw neat diagrams of a CNN pipeline and an LSTM cell.
- These are practice questions, not past papers; check your own syllabus and official previous papers.

About these important questions

This page lists the most important Deep Learning questions for B.Tech CSE and IT exams, arranged unit by unit. The units follow the standard syllabus that most Indian technical universities share, so the order may differ slightly from yours.

Please note: these are practice questions based on common syllabus topics. They are not verified past papers, and we do not claim that any question will appear in your exam. Check your own university syllabus first, then solve the official previous year papers. AKTU students can download them from the AKTU old question paper archive.

Unit 1: Neural Network Basics

Covers the biological neuron, perceptrons, multilayer networks and activation functions.

Short questions (2 marks)

- Differentiate between machine learning and deep learning.
- Why can a single perceptron not solve XOR?
- List four activation functions.

Long questions

1. Explain the perceptron model and the perceptron learning rule, and train it for the AND gate.
2. Explain a multilayer perceptron and the universal approximation idea.
3. Compare sigmoid, tanh, ReLU, leaky ReLU and softmax with their ranges and uses.
4. Explain forward propagation with a small numeric example.
5. Explain loss functions: mean squared error and cross-entropy.

Unit 2: Training Deep Networks

Covers backpropagation, optimisation and the problems of training deep models.

Short questions (2 marks)

- What is the vanishing gradient problem?
- Differentiate between batch, mini-batch and stochastic gradient descent.
- What is an epoch?

Long questions

1. Derive the backpropagation algorithm for a network with one hidden layer.
2. Explain the vanishing and exploding gradient problems and their solutions.
3. Explain momentum, RMSProp and Adam optimisers.
4. Explain weight initialisation methods (Xavier and He).
5. Explain learning rate schedules and early stopping.

Unit 3: Regularisation and Convolutional Neural Networks

Covers regularisation and CNNs. This unit often carries the most marks.

Short questions (2 marks)

- What is dropout?
- What is padding in a CNN?
- What does a pooling layer do?

Long questions

1. Explain L2 regularisation, dropout, batch normalisation and data augmentation.
2. Explain the architecture of a CNN: convolution, activation, pooling and fully connected layers.
3. For a 32 × 32 × 3 input, 10 filters of 5 × 5, stride 1 and no padding, find the output size and the number of parameters. *Hint:* output = (W − F + 2P) / S + 1 = 28, so 28 × 28 × 10; parameters = (5 × 5 × 3 + 1) × 10 = 760.
4. Explain LeNet, AlexNet, VGG and ResNet, and the idea of skip connections.
5. Explain transfer learning and fine-tuning.

Unit 4: Sequence Models

Covers RNNs, LSTM, GRU and attention.

Short questions (2 marks)

- Why are RNNs used for sequence data?
- List the gates in an LSTM.
- What is backpropagation through time?

Long questions

1. Explain the RNN architecture with an unrolled diagram, and its types (one-to-many, many-to-one, many-to-many).
2. Explain the LSTM cell with a diagram and the role of each gate.
3. Compare LSTM and GRU.
4. Explain bidirectional RNNs and encoder-decoder (sequence-to-sequence) models.
5. Explain the attention mechanism and the basic idea of the transformer.

Unit 5: Autoencoders, Generative Models and Applications

Covers autoencoders, GANs and applications. Some universities add deep reinforcement learning here.

Short questions (2 marks)

- What is an autoencoder?
- What are the generator and discriminator in a GAN?

Long questions

1. Explain autoencoders and their types: undercomplete, denoising, sparse and variational.
2. Explain the architecture and training of a GAN, and the problem of mode collapse.
3. Explain applications of deep learning in computer vision, speech and NLP.
4. Explain word embeddings (Word2Vec) in brief.
5. Explain deep reinforcement learning with the example of a deep Q-network.

Sample answer outline: Backpropagation

Backpropagation is the most repeated long question. A full answer has six parts:

1. **State the goal.** Find how much each weight contributes to the error, so gradient descent can update it: $w = w - \eta \times \partial E/\partial w$.
2. **Draw the network.** Input layer, one hidden layer and an output layer, with weights labelled.
3. **Forward pass.** Compute the net input and activation for each hidden unit, then for the output unit. Write the error, for example $E = \frac{1}{2} (t - o)^2$.
4. **Output layer gradient.** With sigmoid activation, $\delta(o) = (t - o) \times o \times (1 - o)$. The update for a hidden-to-output weight is $\Delta w = \eta \times \delta(o) \times h$.
5. **Hidden layer gradient.** $\delta(h) = h \times (1 - h) \times$ sum of $(w \times \delta(o))$ over the outputs it connects to. The update for an input-to-hidden weight is $\Delta w = \eta \times \delta(h) \times x$.
6. **Repeat.** Update all weights, then repeat for every example and epoch until the error is small.

Close with one line on the chain rule, which is what makes backpropagation work, and mention the vanishing gradient problem with sigmoid.

Marks weightage and how to write answers

A common pattern, used by AKTU and many other universities, has a compulsory short-answer section that covers every unit, followed by long-answer questions with a choice inside each unit. Marks per question change between universities and years, so confirm the scheme on a recent official paper.

- **Short answers (about 2 marks):** a crisp definition, one example or diagram, and nothing more. Three to five lines is enough.
- **Long answers (7 to 10 marks):** start with a definition, add a labelled diagram or table, explain the steps, and finish with an example, advantages or a comparison.
- Draw diagrams for every architecture question: a CNN pipeline, an LSTM cell, an autoencoder or a GAN.
- For numericals such as CNN output size, write the formula, substitute, and state the units (height × width × depth).

A 7-day revision plan for Deep Learning

1. **Days 1 to 5:** one unit a day. Read your notes, then write the long answers for that unit from memory and check them against your notes.
2. **Day 6:** redo the backpropagation derivation, two CNN output size problems, and draw the LSTM and GAN diagrams from memory.
3. **Day 7:** solve one full official previous year paper in three hours, then check which units cost you time.

Keep a single sheet of formulas, definitions and diagrams for the last hour before the exam.

Frequently asked questions

Which Deep Learning topics are most important?

Perceptrons, activation functions, backpropagation, optimisers, regularisation, CNN architecture and output size, RNN and LSTM, autoencoders and GANs are asked most often.

Do I need advanced maths for Deep Learning exams?

You need basic calculus (derivatives and the chain rule), matrix multiplication and probability. Practise the backpropagation derivation until the steps feel natural.

How can I practise Deep Learning beyond the exam?

Build small projects with a framework such as PyTorch or TensorFlow, for example an image classifier with a CNN or a sentiment model with an LSTM.

Are these Deep Learning questions from previous year papers?

No. They are practice questions built from topics that appear in most B.Tech Deep Learning syllabi. Use them with your own university's official previous papers, for example the AKTU archive at https://aktu.ac.in/old-question-paper.html.

Is this PDF free?

Yes. Fill in the short form on this page and the PDF downloads right away.

Sources and further reading

- AKTU: old question papers (official archive)
- Deep Learning (Goodfellow, Bengio and Courville), free online book
- NPTEL: online courses from IITs (search for Deep Learning)

About this PDF

These are practice questions for revision, prepared by the Uncodemy trainer team from common B.Tech syllabus topics. They are not past papers, a certificate or a promise of marks. Spot an error? Write to info@uncodemy.com and the team will review it.

More free resources and courses: uncodemy.com | info@uncodemy.com | +91 98183 66550